TRACKING FOR THE ATLAS LEVEL 1 TRIGGER FOR THE HIGH LUMINOSITY LHC

M. R. Sutton[†]

on behalf of the ATLAS Collaboration

*The Department of Physics and Astronomy,
The University of Sussex,*

Pevensey II, University of Sussex, Falmer, Brighton BN1 9QH, Great Britain

[†]*E-mail: sutt@cern.ch*

At the HL-LHC, the increased luminosity will result in up to 200 pile-up interactions per bunch crossing. One of the greatest challenges for ATLAS will be to keep the Level 1 Trigger p_T thresholds low enough to maintain high trigger efficiency for all interesting physics. The proposed two-stage design of the ATLAS Level 1 Trigger, and the incorporation of a Level-1 track trigger is described. The requirements and implications for the tracker readout architecture, and estimates of readout latency based on a detailed discrete event simulation of the data flow in the tracker front-end electronics are also presented.

Keywords: Trigger, Tracking, ATLAS, Level-1, LHC

1. Introduction

During the 2012 data taking period the LHC¹ performed extremely well, delivering more than 20 fb^{-1} for both the ATLAS² and CMS³ experiments with mean numbers of interactions per bunch crossings of more than 20. Even at these high luminosities, many important physics channels are, and will continue to be for the foreseeable future, statistics limited.

In order to help overcome this limitation and allow increased phase space coverage for searches and precision Higgs physics, a luminosity upgrade for the LHC is foreseen for 2024. For this high luminosity LHC (HL-LHC), instantaneous luminosities around $7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ are expected, delivering around 300 fb^{-1} per year.⁴ At these high luminosities, mean numbers of separate pp interactions in the range of 160-200 are expected for every bunch crossing. Event rates for single lepton candidates passing the Level 1 (L1) triggers in the ATLAS detector using the current trigger strategy are expected to be of the order of 200-400 kHz for *each* lepton species individually, with jet and other physics signatures providing some significant rate in addition. The current ATLAS trigger strategy⁵ uses only coarse granularity Calorimeter and Muon Spectrometer information at Level 1, with full granularity data including tracking information from the Inner Detector used

in the High Level Trigger (HLT). Processing this rate of output from L1 is not possible in the HLT.

2. Tracking for the ATLAS Level 1 trigger

For the HL-LHC upgrade ATLAS will be equipped with an all-new Silicon Inner Tracker (ITK), consisting of 6 layers of Si strip detectors in the barrel and 7 in the endcaps, enclosing 4 layers of high precision pixel detectors in the barrel and 6 in the endcaps.

Studies have been performed⁶ to evaluate the performance benefits of reconstructing tracks in the ITK at Level 1. By matching tracks to Calorimeter clusters used to identify electrons or taus, or, to tracks from the Level 1 Muon Spectrometer reconstruction for muons, it will be possible to reduce the rate of electron or tau candidates at Level 1 by factors of between three and ten, and reduce the rate of muon candidates by factors approaching three, with only a small (<10%) loss of efficiency. These levels of rejection for each lepton trigger are possible even with reasonably modest tracking performance at Level 1.

2.1. The split Level 0 – Level 1 scenario

To enable the processing of the high rates from the HL-LHC, ATLAS has therefore proposed a Level 1 Track Trigger to work within a scenario where the current Level 1 processing is split into consecutive Level 0 (L0) and Level 1 hardware stages.⁶ This is illustrated in Figure 1.

Limitations from the front-end pipelines for no-longer accessible sub-detector components mean that any upgraded ATLAS trigger system will have only 20 μ s from the time of the bunch crossing until the L1 decision must arrive back at the detector front ends. The maximum L1 output rate that can be supported by these sub-detectors is 200 kHz.

To operate within these constraints, in the split Level 0 – Level 1 scenario the detector front end electronics for some sub-detectors – such as the ITK – will make use of a *double buffer*; the first, *primary* buffer clocked at the bunch crossing rate, as in the current L1 system. In the L0 system the Calorimeter and Muon systems will process events using coarse granularity data to identify regions of interest (RoI) as in the current Level 1 trigger. The L0 decision whether to accept or reject each event will be distributed back to the detector front ends at the rate of 500 kHz to arrive 6 μ s after each bunch crossing. On a L0 accept the data from the primary buffers will be read into a *secondary* buffer where it will remain until the L1 decision

Fig. 1. The proposed split Level 0/Level 1 architecture for the ATLAS Phase II upgrade⁶ showing the different L0 sub-systems for the Calorimeter and Muon Spectrometer.

arrives after a further $14\mu\text{s}$. Upon a L1 reject decision the data is simply dropped. Upon a L1 accept (L1A) it is read out from the secondary buffer into the ATLAS off-detector DAQ system. Within this $14\mu\text{s}$ both the transfer of the ITK data to the L1 track processor and the track reconstruction must be performed. The readout latency is critical – if the data can be read out within $6\mu\text{s}$ then $8\mu\text{s}$ will remain for the tracking algorithm.

In order to integrate the L1 Track Trigger, the location of any RoIs identified by the Calorimeter or Muon systems will be distributed back to the detector front ends via the Readout Drivers (RODs) in the form of a *Regional Readout Request* (R3). Upon receipt of an R3, each detector element within the RoI will read out a sub-sample of the data from the secondary buffer to the L1 processing systems. For each L0 accept it is expected that the average number and size of all the RoIs in each event will comprise only about 10% of the total data volume within the event. In this way, only 10% of the detector need be read out for processing by L1 on any L0 accept, such that the total event rate to be read out to the L1 system from any single part of the detector will be 50 kHz – only 10% of the 500 kHz L0 accept rate.

3. The ATLAS Inner Tracker strip readout

The readout of the ITK strip tracker is arranged by module – each module reading out the strips using several daisy-chained, identical front end ASICs

Fig. 2. A schematic of the functionality of the ABC130 FE chip.

(ABC130) which use 130 nm CMOS technology, and which are connected to a Hybrid Chip Controller (HCC). The chips connected to each HCC are arranged into two separate daisy chains. Each chain can be addressed from both directions for redundancy should any chip within the chain fail. In the barrel, there are 10 chips per hybrid, read out by each HCC, arranged into 2 chains of 5 chips. For the endcaps, the number of chips attached to each HCC depends on the location of the module. For these studies 12 chips are connected to the endcap HCC with the most chips.

Data packets from the HCC for all the modules on each stave (or each petal in endcaps) are aggregated and transferred from the stave using a GBT link allowing data to be sent from each Hybrid at a rate of 160 Mbps.

Figure 2 illustrates the functionality of the ABC130 readout. This contains logic to prioritise the read out of R3 data over that of data from a L1 accept. Following an R3, after generating and reading out any R3 data packet encoding the cluster information for the chip, a *duty cycle* algorithm will prioritise any packets being sent from the previous chips in the daisy chain over any subsequent R3 or L1 accept request. In this way, all the packets from any given R3 should arrive more quickly at the HCC.

3.1. Emulating the ITK strip tracker readout

Studies using a discrete event simulation of the complete readout chain from the distribution of the R3 and L1 accept requests, to readout through

Fig. 3. The latency to read-out of all the packets for 95% of all R3 requests for all the hybrids in the innermost barrel layer of the ITK strip tracker from the discrete event simulation. Shown are the times with (solid) and without (dashed) inclusion of prioritisation of R3 data sending on the HCC.

the HCC to the off-detector electronics have been performed using ITK hit occupancies expected with 200 pileup interactions per bunch crossing. These suggest that the “R3 95% latency” – the time taken to read out all packets from a hybrid for 95% of all R3 – for all modules in the inner most barrel layer of the ITK strip tracker, is safely within $6\ \mu\text{s}$ for the standard 200 kHz L1A rate. This is illustrated in Figure 3 which shows the R3 95% latency for all the hybrids in the highest occupancy barrel stave, closest to the beamline, as a function of the L1A rate. Shown are the latencies both with, and without R3 data prioritisation on the HCC. Without the HCC prioritisation, R3 data are queued behind earlier L1A data. Shorter readout latencies are expected for barrel modules with lower occupancies.

For the endcap, the modules are less homogeneous, with different chip multiplicities per hybrid. Hybrid rings 4 - 6 have 11, 11 and 12 chips respectively attached to the HCC. The larger number of chips increases both the data volume, and the length of the daisy chains, so increasing the time and number of steps required to transfer packets from the last chip in the chain to the HCC. In simulation, hybrid ring 6 is most problematic taking well over $6\ \mu\text{s}$, even at low L1A rates. This is illustrated in Figure 4. In order to prioritise the R3 data on the HCC and prevent data blocking on the daisy chains, an additional FIFO has been added to buffer the input on the HCC from each daisy chain. Latencies are shown for the cases of a FIFO with unlimited depth, and with a maximum depth of 32 packets.

For these problematic hybrids, the discrete event simulation suggests that making use of additional GBT bandwidth available for the endcap petals, and using all four readout links from the daisy-chains on the hybrid to ship the data to the HCC more quickly should allow all packets to be read out in under $6\ \mu\text{s}$ and all the packets from 95% of all R3 within $5\ \mu\text{s}$.

Fig. 4. The latency to read out all the packets for 95% of all R3 requests for all hybrids (labeled Hyb0-6) from the most forward endcap wheel in the ITK strip tracker from the discrete event simulation. Shown are the times with an unlimited, and 32 packet deep FIFO on the HCC inputs from the chip daisy chains.

4. Outlook

Using the current ABC130 design and implementing R3 prioritisation on the HCC it should be possible to readout the entire ITK strip tracker barrel in under $6\mu\text{s}$ for 95% of all regional readout requests. Making use of spare bandwidth available in the endcaps and redundancy in the hybrid daisy chain links to the HCC it should be possible to read out the entire endcap also in under $6\mu\text{s}$, to leave around $8\mu\text{s}$ for the track reconstruction.

The prototype ABC130 chip design has been submitted for fabrication and chips should be available for testing and development soon. The HCC design has not advanced as far the ABC130 with the current design concentrating on the functionality to control the ABC130 daisy chains. However the next version of the HCC design, foreseen for fabrication in 2015 will include the provision to prioritise the R3 data.

References

1. O. Brüning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, J. Poole and P. Proudlock, *LHC Design Report* (CERN, Geneva, 2004).
2. The ATLAS Collaboration, *Journal of Instrumentation* **3**, p. S08003 (2008).
3. The CMS Collaboration, *Journal of Instrumentation* **3**, p. S08004 (2008).
4. O. Brüning, L. Rossi, *et al*, *High-Luminosity Large Hadron Collider; A description for the European Strategy Preparatory Group*, Tech. Rep. CERN-ATS-2012-236, CERN (Geneva, 2012).
5. The ATLAS Collaboration, *Eur.Phys.J.* **C72**, p. 1849 (2012).
6. The ATLAS Collaboration, *Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment*, Tech. Rep. CERN-LHCC-2012-022 ; LHCC-I-023, CERN (Geneva, 2013).